

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	23	migrat\$5 same stream\$4 same data same stor\$4 same locati\$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/31 12:42
L2	3	L1 and (portion\$4 same relocat\$5 same (metadata meta-data) same access\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/31 12:48
L3	0	L2 and (preserv\$4 same stream\$4 same relationship\$1 same (meta-data metadata) same access\$4 same migrat\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/31 12:50
L4	1	(preserv\$4 same stream\$4 same relationship\$1 same (meta-data metadata) same access\$4 same migrat\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/31 12:50



[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

(preserv\$4 same stream\$4 same relationship\$1 same (meta-d



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

preserv\$4 same stream\$4 same relationship\$1 same meta data metadata same access\$4 same migrat\$4

11
11

Sort results by

[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

[Search Tips](#)

[Try this search in The ACM Guide](#)

☐ [Open results in a new window](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐

1 [The measured performance of personal computer operating systems](#)

J. B. Chen, Y. Endo, K. Chan, D. Mazieres, A. Dias, M. Seltzer, M. D. Smith

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: [pdf\(1.98 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

2 [Improving storage system availability with D-GRAID](#)

Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Full text available: [pdf\(700.30 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of D-GRAID, a gracefully degrading and quickly recovering RAID storage array. D-GRAID ensures that most files within the file system remain available even when an unexpectedly high number of faults occur. D-GRAID achieves high availability through aggressive replication of semantically critical data, and fault-isolated placement of logically related data. D-GRAID also recovers from failures quickly, restoring only live file system data to a h ...

Keywords: Block-based storage, Disk array, RAID, fault isolation, file systems, smart disks

3 [Querying web metadata: Native score management and text support in databases](#)

Gültekin Özsoyoglu, Ismail Sengör Altıngövdé, Abdullah Al-Hamdani, Selma Ayşe Özel, Özgür Ulusoy, Zehra Meral özsoyoglu

December 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 4

Full text available: [pdf\(737.76 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


In this article, we discuss the issues involved in adding a native score management system to object-relational databases, to be used in querying Web metadata (that describes the semantic content of Web resources). The Web metadata model is based on topics (representing entities), relationships among topics (called *metalinks*), and importance scores (sideway values) of topics and metalinks. We extend database relations with scoring functions and importance scores. We add to SQL score-manag ...

Keywords: Score management for Web applications

4 Practical byzantine fault tolerance and proactive recovery

Miguel Castro, Barbara Liskov

November 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 4

Full text available:  pdf(1.63 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

5 File and storage systems: The Google file system

Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available:  pdf(275.54 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: clustered storage, data storage, fault tolerance, scalability

6 Operating systems: DualFS: a new journaling file system without meta-data duplication

Juan Piernas, Toni Cortes, José M. García

June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available:  pdf(213.64 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we introduce DualFS, a new high performance journaling file system that puts data and meta-data on different devices (usually, two partitions on the same disk or on different disks) and manages them in very different ways. Unlike other journaling file systems, DualFS has only one copy of every meta-data block. This copy is in the *meta-data device*, a log which is used by DualFS both to read and to write meta-data blocks. By avoiding a time-expensive extra copy of meta-data b ...

Keywords: DualFS, journaling file system, meta-data management

7 Soft updates: a solution to the metadata update problem in file systems

Gregory R. Ganger, Marshall Kirk McKusick, Craig A. N. Soules, Yale N. Patt

May 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 2

Full text available:  pdf(147.90 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Metadata updates, such as file creation and block allocation, have consistently been identified as a source of performance, integrity, security, and availability problems for file systems. Soft updates is an implementation technique for low-cost sequencing of fine-grained updates to write-back cache blocks. Using soft updates to track and enforce metadata update dependencies, a file system can safely use delayed writes for almost all file operations. This article describes soft ...

8 Dynamic Metadata Management for Petabyte-Scale File Systems

Sage A. Weil, Kristal T. Pollack, Scott A. Brandt, Ethan L. Miller

November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Full text available:  pdf(175.04 KB)

Additional Information: [full citation](#), [abstract](#)

In petabyte-scale distributed file systems that decouple read and write from metadata operations behavior of the metadata server cluster will be critical to overall system performance and scalability. We present a dynamic subtree partitioning and adaptive metadata management system designed to efficiently manage hierarchical metadata workloads that evolve over time. We examine the relative merits of our approach in the context of traditional workload partitioning strategies, and demonstrate the ...

9 **BASE: Using abstraction to improve fault tolerance**

Miguel Castro, Rodrigo Rodrigues, Barbara Liskov

August 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 3

Full text available:  pdf(438.18 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability ...

Keywords: Byzantine fault tolerance, N-version programming, asynchronous systems, proactive recovery, state machine replication

10 **Industrial papers: metadata management for data integration: Meta-data version and configuration management in multi-vendor environments**

John R Friedrich

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(544.92 KB)


Additional Information: [full citation](#), [abstract](#), [references](#)

Nearly all components that comprise modern information technology, such as Computer Aided Software Engineering (CASE) tools, Enterprise Application Integration (EAI) environments, Extract/Transform/Load (ETL) engines. Warehouses, EII, and Business Intelligence (BI), contain a great deal of meta-data, which often drive much of the tool's functionality. These metadata are distributed and duplicated, are often times actively interacting with the tools as they process data and are generally represe ...

11 **ViSWeb—the Visual Semantic Web: unifying human and machine knowledge representation with Object-Process Methodology**

Dov Dori

May 2004 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 1 Issue 2

Full text available:  pdf(1.22 MB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

The Visual Semantic Web (ViSWeb) is a new paradigm for enhancing the current Semantic Web technology. Based on Object-Process Methodology (OPM), which enables modeling of systems in a single graphic and textual model, ViSWeb provides for representation of knowledge over the Web in a unified way that caters to human perceptions while also being machine processable. The advantages of the ViSWeb approach include equivalent graphic-text knowledge representation, visual navigability, semantic sentenc ...

Keywords: Conceptual graphs, Knowledge representation, Object-Process Methodology, Semantic Web, Visual Semantic Web

12 DLFM: a transactional resource manager

Hui-I Hsiao, Inderpal Narang

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2

Full text available:  [pdf\(124.99 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The DataLinks technology developed at IBM Almaden Research Center and now available in DB2 UDB 5.2 introduces a new data type called DATALINK for a database to reference and manage file stored external to the database. An external file is put under a database control by "linking" the file to the database. Control to a file can also be removed by "unlinking" it. The technology provides transactional semantics with respect to linking or unlinking the file when DATALINK ...

13 Applications: Web-assisted annotation, semantic indexing and search of television and radio news

Mike Dowman, Valentin Tablan, Hamish Cunningham, Borislav Popov

May 2005 **Proceedings of the 14th international conference on World Wide Web**

Full text available:  [pdf\(403.97 KB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Rich News system, that can automatically annotate radio and television news with the aid of resources retrieved from the World Wide Web, is described. Automatic speech recognition gives a temporally precise but conceptually inaccurate annotation model. Information extraction from related web news sites gives the opposite: conceptual accuracy but no temporal data. Our approach combines the two for temporally accurate conceptual semantic annotation of broadcast news. First low quality transcri ...

Keywords: Web search, automatic speech recognition, key-phrase extraction, media archiving, multimedia, natural language processing, semantic Web, semantic annotation, topical segmentation

14 Workshop on compositional software architectures: workshop report

May 1998 **ACM SIGSOFT Software Engineering Notes**, Volume 23 Issue 3


Full text available:  [pdf\(2.91 MB\)](#)

Additional Information: [full citation](#), [index terms](#)

15 Application performance and flexibility on exokernel systems

M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector M. Briceño, Russell Hunt, David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles**, Volume 31 Issue 5

Full text available:  [pdf\(2.39 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

16 Negotiation-based protocols for disseminating information in wireless sensor networks

Joanna Kulik, Wendi Heinzelman, Hari Balakrishnan

March 2002 **Wireless Networks**, Volume 8 Issue 2/3

Full text available:  [pdf\(357.49 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we present a family of adaptive protocols, called SPIN (Sensor Protocols for Information via Negotiation), that efficiently disseminate information among sensors in an energy constrained wireless sensor network. Nodes running a SPIN communication protocol name their data using high-level data descriptors, called meta-data. They use meta-data negotiations to


eliminate the transmission of redundant data throughout the network. In addition, SPIN nodes can base their communication decisions on ...

Keywords: energy-efficient protocols, information dissemination, meta-data, negotiation-based protocols, wireless sensor networks

17 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available:  [pdf\(4.21 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial communication ...

18 Using semantic values to facilitate interoperability among heterogeneous information systems

Edward Sciore, Michael Siegel, Arnon Rosenthal

June 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 2

Full text available:  [pdf\(2.68 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Large organizations need to exchange information among many separately developed systems. In order for this exchange to be useful, the individual systems must agree on the meaning of their exchanged data. That is, the organization must ensure semantic interoperability. This paper provides a theory of semantic values as a unit of exchange that facilitates semantic interoperability between heterogeneous information systems. We show how semantic values can be used to ...

19 BASE: using abstraction to improve fault tolerance

Rodrigo Rodrigues, Miguel Castro, Barbara Liskov

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles**, Volume 35 Issue 5

Full text available:  [pdf\(1.47 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability ...

20 Ontology-based metadata: transforming the MARC legacy

Peter C. Weinstein

May 1998 **Proceedings of the third ACM conference on Digital libraries**

Full text available:  [pdf\(1.21 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)